

# LogitBoost autoregressive networks

Marc Goessling\*

*Department of Statistics, University of Chicago, USA*

---

## Abstract

Multivariate binary distributions can be decomposed into products of univariate conditional distributions. Recently popular approaches have modeled these conditionals through neural networks with sophisticated weight-sharing structures. It is shown that state-of-the-art performance on several standard benchmark datasets can actually be achieved by training separate probability estimators for each dimension. In that case, model training can be trivially parallelized over data dimensions. On the other hand, complexity control has to be performed for each learned conditional distribution. Three possible methods are considered and experimentally compared. The estimator that is employed for each conditional is LogitBoost. Similarities and differences between the proposed approach and autoregressive models based on neural networks are discussed in detail.

*Keywords:* Autoregressive model, Conditional distribution, Boosting

---

## 1. Introduction

Given a collection of multivariate binary data points  $\mathbf{x}^{(n)} \in \{0,1\}^D$ ,  $n = 1, \dots, N$ , we consider the fundamental task of learning the underlying distribution  $\mathbb{P}(\mathbf{x})$ . Distribution estimates have many potential applications. For example, they can be used for anomaly detection, denoising, data synthesis, missing value imputation, classification and compression. A useful fact is that every multivariate distribution can be factored into a product of univariate conditional distributions

$$\mathbb{P}(\mathbf{x}) = \prod_{d=1}^D \mathbb{P}(x_d | \mathbf{x}_{1:d-1}). \quad (1)$$

Thus, it is always possible to learn a joint distribution by estimating univariate conditionals. The decomposition (1) is known as a fully connected left-to-right graphical model or autoregressive network (Frey, 1998). Fitting a sequence of univariate conditional distributions can be easier than trying to model high-dimensional dependencies directly. Moreover, with the factorization it is straightforward to perform likelihood evaluations and exact sampling.

Many different approaches have been proposed to model the conditional distributions in an autoregressive network. In the simplest approach (Cooper and Herskovits, 1992) the conditional distributions are learned by performing a greedy search for a sparse set  $\text{par}(d) \subset \{1, \dots, d-1\}$  of predictors for each dimension  $d$ . The conditional probabilities  $P(x_d | \mathbf{x}_{\text{par}(d)})$  are then simply estimated through the corresponding empirical frequencies. However, since the number of parameters that have to be learned grows exponentially in the number of predictors, only very few predictors can be used.

A fruitful alternative is to approximate the conditional distributions through parametric models. For example, Frey (1998) considered logistic autoregressive networks, which are also called fully visible sigmoid

---

\*Corresponding author.

*Email address:* goessling@uchicago.edu (Marc Goessling)

belief nets. In these networks the log-odds of  $x_d = 1$  are modeled as a linear function of  $\mathbf{x}_{1:d-1}$ . That means

$$\mathbb{P}(x_d=1 \mid \mathbf{x}_{1:d-1}) = \sigma(\alpha_1^{(d)}x_1 + \dots + \alpha_{d-1}^{(d)}x_{d-1}),$$

where  $\boldsymbol{\alpha}^{(d)} \in \mathbb{R}^{d-1}$  and  $\sigma(z) = (1 + \exp(-z))^{-1}$  is the logistic function. However, the performance of these simple networks can be limited if the true conditional distributions are highly nonlinear. Mixtures of such networks (Goessling and Amit, 2016) often lead to improvements if sufficient data is available.

More ambitious approaches attempt to directly model the log-odds of  $x_d = 1$  as a nonlinear function of  $\mathbf{x}_{1:d-1}$ . One major research direction focuses on neural networks for this task. These models are of the form

$$\mathbb{P}(x_d=1 \mid \mathbf{x}_{1:d-1}) = \sigma(b^{(d)} + (\mathbf{v}^{(d)})^T \mathbf{h}^{(d)}), \quad (2)$$

where  $b^{(d)} \in \mathbb{R}$ ,  $\mathbf{v}^{(d)} \in \mathbb{R}^H$ . In the original work by Bengio and Bengio (2000), each data dimension was modeled through a different set of (deterministic) hidden units

$$\mathbf{h}^{(d)} = \sigma(\mathbf{c}^{(d)} + \mathbf{W}^{(d)}\mathbf{x}_{1:d-1}) \in \mathbb{R}^H,$$

where  $\mathbf{c}^{(d)} \in \mathbb{R}^H$ ,  $\mathbf{W}^{(d)} \in \mathbb{R}^{H \times (d-1)}$ . However, that approach suffered from substantial overfitting. To improve the generalization performance it was suggested to prune connections after the training phase, based on statistical tests. In other words, some entries in  $\mathbf{W}^{(d)}$  were manually set to 0. Larochelle and Murray (2011) subsequently introduced the idea of weight sharing between the parameters for different dimensions. Specifically, they proposed to use the same bias term  $\mathbf{c}^{(d)} = \mathbf{c}$  for all dimensions  $d$  and connections of the form  $\mathbf{W}^{(d)} = \mathbf{W}_{:,1:d-1}$ , where  $\mathbf{W} \in \mathbb{R}^{H \times (D-1)}$ . This approach led to drastically better results (Bengio, 2011). Many variations of their idea (e.g. Uria et al., 2014; Raiko et al., 2014; Germain et al., 2015) have been considered afterwards. The main difference between the proposed models is the number of hidden layers, the connections within and between layers, and the choice of the activation function. The more advanced models use an ensemble of networks corresponding to different orderings of the data dimensions or corresponding to different connectivity patterns. An equal mixture of these networks is then used as the final model.

A rather different class of nonparametric estimators are decision trees (Breiman et al., 1984). When employed for probabilistic modeling they are also known as probability estimation trees (Provost and Domingos, 2003). These trees model conditional probabilities through piecewise-constant functions

$$\mathbb{P}(x_d=1 \mid \mathbf{x}_{1:d-1}) = \sum_{j=1}^J p_j^{(d)} \mathbb{1}(\mathbf{x}_{1:d-1} \in R_j^{(d)}),$$

where the regions  $R_j^{(d)}$  partition the predictor space  $\{0,1\}^{d-1}$  and  $p_j^{(d)} \in (0,1)$ . Since the predictors are binary, each region is defined by specifying the value of some of the variables  $x_1, \dots, x_{d-1}$ . Decision trees as models for conditional distributions have been used for example by Boutilier et al. (1996); Friedman and Goldszmidt (1998). Decision trees are attractive because they are flexible, interpretable and computationally efficient. The performance of single trees is often unsatisfactory though because shallow trees can have a large bias and deep trees often have a large variance.

However, decision trees can be assembled into a powerful estimator through a method called boosting (Friedman, 2001). This procedure creates a collection of relatively shallow decision trees, which are trained as to complement each other. The bias-variance tradeoff for boosted trees is often much better than for single decision trees. Particularly relevant for us is LogitBoost (Friedman et al., 2000), which is a probability estimator based on boosted trees.

In this work, we explore the use of LogitBoost for learning high-dimensional binary distributions from data. We call our model a LogitBoost autoregressive network (LBARN). Our contributions are as follows:

1. We propose LogitBoost as a learning procedure for the conditionals in an autoregressive network. In existing work (Shafik and Tutz, 2009) boosting was applied to continuous, stationary time series. Thus, only a single conditional distribution had to be learned. Our network, on the other hand, can train several hundred conditional distributions (one for each data dimension).

2. Neural autoregressive networks are currently the state-of-the-art for nonparametric distribution estimation in high-dimensions (say  $D > 50$ ). We show that there is a simpler alternative that works equally well. Similarities and differences between our approach and neural methods are discussed in detail.
3. In contrast to (most) neural autoregressive networks our model does not use any weight sharing among conditionals for different dimensions. Consequently, our training procedure can be perfectly parallelized over the data dimensions. However, this also means that we need a method to perform complexity control for each conditional distribution. We propose three different procedures and empirically evaluate them on several diverse datasets.
4. Finally, we study how the ordering of the variables affects the performance of the trained model. In particular, we introduce a sorting procedure that allows us to build a simple compression mechanism.

We start in Section 2 by providing an overview of LogitBoost. In Section 3 we then describe how separate LogitBoost runs for each dimension can be combined into a coherent joint model. In Section 4 we contrast our model with neural autoregressive networks. In Section 5 we present quantitative results for several common datasets and discuss the choice of hyperparameters. We also show how the strength of (nonlinear) dependencies between the variables can be quantified in our model and illustrate that the conditional distributions can be learned robustly even when a large number of irrelevant variables are present. Moreover, we consider how the autoregressive network is affected when different learned orderings of the variables are used.

## 2. LogitBoost

LogitBoost (Friedman et al., 2000) is a forward stagewise procedure that fits additive logistic regression models<sup>1</sup>

$$\mathbb{P}_t(y=1 \mid \mathbf{x}) = \sigma(f_1(\mathbf{x}) + \dots + f_t(\mathbf{x}))$$

for  $t = 1, \dots, T$  to the training data  $(\mathbf{x}^{(n)}, y^{(n)}) \in \{0, 1\}^D \times \{0, 1\}$ ,  $n = 1, \dots, N$ , by maximum likelihood. While in the context of generalized additive models (Hastie and Tibshirani, 1990) additivity refers to a sum of (smooth) univariate functions, here additivity is meant in the sense of a sum of multivariate functions. The optimization is based on a second-order Taylor expansion of the Bernoulli log-likelihood

$$L_t = \sum_{n=1}^N \log \mathbb{P}_t(y^{(n)} \mid \mathbf{x}^{(n)}).$$

Friedman et al. (2000) derived that

$$\frac{\partial L_t}{\partial f_t} = \sum_{n=1}^N (y^{(n)} - p_{t-1}^{(n)}), \quad \frac{\partial^2 L_t}{\partial f_t^2} = - \sum_{n=1}^N p_{t-1}^{(n)} (1 - p_{t-1}^{(n)}),$$

where

$$p_{t-1}^{(n)} = \mathbb{P}_{t-1}(y=1 \mid \mathbf{x}^{(n)})$$

is the conditional success probability given  $\mathbf{x}^{(n)}$  as predicted by the model after  $t - 1$  rounds of boosting. The Newton step results in a weighted least-squares regression of the current pseudoresiduals  $y^{(n)} - p_{t-1}^{(n)}$  onto the predictors  $\mathbf{x}^{(n)}$ . The base learners are typically  $J$ -terminal regression trees

$$f_t(\mathbf{x}) = \sum_{j=1}^J \gamma_j \mathbb{1}\{\mathbf{x} \in R_j\},$$

---

<sup>1</sup>The model  $\mathbb{P}_0$  is defined as a Bernoulli-1/2 distribution, independently of  $\mathbf{x}$ .

where  $\gamma_j \in \mathbb{R}$ . It can be shown (Li, 2010) that the weighted regression problem is then equivalent to fitting a decision tree with a simple objective function. Specifically, the objective value for a subset  $R$  of the predictor space  $\{0, 1\}^D$  is

$$\text{obj}(R) = \frac{\left[ \sum_{n=1}^N (y^{(n)} - p_{t-1}^{(n)}) \mathbb{1}(\mathbf{x}^{(n)} \in R) \right]^2}{\sum_{n=1}^N p_{t-1}^{(n)} (1 - p_{t-1}^{(n)}) \mathbb{1}(\mathbf{x}^{(n)} \in R)}.$$

The task is to choose the regions  $R_j \subset \{0, 1\}^D$ ,  $j = 1, \dots, J$ , in such a way that  $\sum_{j=1}^J \text{obj}(R_j)$  is maximized. Since the exact optimization is very expensive, this is usually done approximately through a greedy procedure (Breiman et al., 1984). One starts with a single region, corresponding to the entire predictor space  $\{0, 1\}^D$ , and then recursively divides the region that leads to the largest improvement of the objective. The gain of splitting a region  $R$  using variable  $x_d$  is

$$\text{gain}(d) = \text{obj}(\{\mathbf{x} \in R \mid x_d=0\}) + \text{obj}(\{\mathbf{x} \in R \mid x_d=1\}) - \text{obj}(R). \quad (3)$$

Once a  $J$ -terminal tree is learned the value of the leaves is set to

$$\gamma_j = \frac{\sum_{n=1}^N (y^{(n)} - p_{t-1}^{(n)}) \mathbb{1}(\mathbf{x}^{(n)} \in R_j)}{\sum_{n=1}^N p_{t-1}^{(n)} (1 - p_{t-1}^{(n)}) \mathbb{1}(\mathbf{x}^{(n)} \in R_j)}. \quad (4)$$

The number of leaves  $J$  impacts the depth of the trees and hence determines the order of interactions that can be modeled. While LogitBoost is rather robust to overfitting for classification tasks, overfitting does occur for probability estimates if too many rounds of boosting are performed (Mease et al., 2007). It is therefore customary to choose the number of trees based on the performance on holdout data. Better results can often be achieved by introducing shrinkage since this facilitates the use of somewhat deeper trees, which otherwise would not generalize well. The log-odds of  $y = 1$  in this case is modeled as  $\nu \sum_{t=1}^T f_t$ , where  $\nu \in (0, 1)$  is a shrinkage parameter. Smaller values of  $\nu$  almost always lead to better generalization performance. However, if  $\nu$  is very small then many trees are needed and hence the computational costs grow accordingly.

### 3. Model selection

Our main proposal is to apply LogitBoost separately for each data dimension  $d = 1, \dots, D$  and to learn the conditional distributions  $\mathbb{P}(x_d \mid \mathbf{x}_{1:d-1})$  using up to  $T$  boosted trees. This yields  $D$  sequences of models  $\mathbb{P}_t(x_d \mid \mathbf{x}_{1:d-1})$ ,  $t = 0, \dots, T$ , with increasing complexity. In order to obtain a single joint model  $\mathbb{P}(\mathbf{x})$  for the full data vector, we have to choose the number of boosting rounds  $t_d \in \{0, \dots, T\}$  for each dimension  $d$ . Thus, in total there are  $(T+1)^D$  candidate models for the joint distribution. In the following we propose three different strategies for selecting one of those candidates.

#### 3.1. Individual selection

A straightforward approach is to perform model validation individually for each dimension  $d$ . This means we compute the likelihood of holdout data  $x_d$  given  $\mathbf{x}_{1:d-1}$  under each of the models  $\mathbb{P}_t(x_d \mid \mathbf{x}_{1:d-1})$ ,  $t = 0, \dots, T$ , and find the number of boosting rounds  $t_d^*$  that corresponds to the model with the highest validation likelihood for the  $d$ -th dimension. The joint distribution is then estimated through

$$\mathbb{P}_{\text{ind}}(\mathbf{x}) = \prod_{d=1}^D \mathbb{P}_{t_d^*}(x_d \mid \mathbf{x}_{1:d-1}).$$

If the validation set is rather small then this approach can actually be expected to overfit the validation data since we are making  $D$  model choices. The next two approaches, on the other hand, perform only a single model selection.

### 3.2. Common selection

A conservative method for model selection is to use the entire data vector  $\mathbf{x}$  for validation and to choose a common number of trees for all dimensions. In this case we compute the validation likelihood for the models  $\prod_{d=1}^D \mathbb{P}_t(x_d | \mathbf{x}_{1:d-1})$ ,  $t = 0, \dots, T$ , and find the maximizer  $t^*$ . The joint distribution is then estimated through

$$\mathbb{P}_{\text{com}}(\mathbf{x}) = \prod_{d=1}^D \mathbb{P}_{t^*}(x_d | \mathbf{x}_{1:d-1}).$$

This approach implicitly assumes that all conditional distributions have about the same complexity. Hence, if the individual variables are heterogeneous (in terms of variance) then the common selection of boosting rounds may hurt the performance of the distribution estimator.

### 3.3. Linearized selection

A third approach for model selection is to construct a linear ordering of all learned trees and to perform validation on the corresponding sequence of joint models. Specifically, we can start with “empty” models  $\mathbb{P}_0(x_d | \mathbf{x}_{1:d-1})$  for all dimensions, i.e., using zero trees. We then greedily add a tree to the dimension that gives rise to the largest improvement in terms of training likelihood for the full data vector  $\mathbf{x}$ . Alternatively, the sorting procedure can be performed in a backward manner by starting with the “full” model  $\mathbb{P}_T(x_d | \mathbf{x}_{1:d-1})$  for each dimension, i.e., using all  $T$  trees. We then greedily remove the tree that gives rise to the smallest decrease of training likelihood. Both procedures create a sequence of models  $\mathbb{P}_s(\mathbf{x})$ ,  $s = 0, \dots, D \cdot T$ . The joint distribution can then be estimated through

$$\mathbb{P}_{\text{lin}}(\mathbf{x}) = \mathbb{P}_{s^*}(\mathbf{x}),$$

where  $s^*$  is the maximizer of the likelihood for holdout data  $\mathbf{x}$ . The linearized selection method makes only a single model selection but allows for different model complexities of the individual dimensions.

## 4. Comparison with neural autoregressive networks

Our model as well as autoregressive models based on neural networks learn a nonlinear function of  $\mathbf{x}_{1:d-1}$  for the (conditional) log-odds of  $x_d = 1$ . We continue with a discussion of similarities and differences between these two approaches.

### 4.1. Parallelization

One major difference is that we estimate separate conditional models for each dimension. This means that our training procedure can be perfectly parallelized over the data dimensions. Indeed, each conditional distribution can be learned on a separate processor without requiring any communication. The fact that state-of-the-art performance can be achieved without any weight sharing (as we will show in Section 5) is a nontrivial result. The parameter sharing introduced by Larochelle and Murray (2011) greatly improved the generalization performance of neural autoregressive networks and there are no reported competitive results for neural approaches without extensive parameter sharing of this type.

### 4.2. Architecture

Another fundamental difference is that neural networks have a fixed architecture, which is chosen by trial and error. This includes the number of hidden layers, the number of units per layer and the connections between them (including the choice of activation function). Only the weights for the connections are learned. In contrast to that, LogitBoost also learns the structure of the nonlinear functions. Indeed, each decision tree produces a function that depends only on a small subset of the variables. The number of tree leaves  $J$  determines what degree of interactions between the predictors can be incorporated. Such a hyperparameter does not exist in neural networks.

#### 4.3. Optimization

A further distinction between the two approaches lies in the optimization procedure. Neural networks use backpropagation, which refers to gradient descent on the log-likelihood of the model. The learning process is started from a random initialization of the parameters. Our model, on the other hand, starts with log-odds of zero for each conditional and uses a second-order Taylor approximation of the log-likelihood for the parameter updates. In both cases the model complexity increases with additional iterations of gradient descent or additional rounds of boosting. The learning rate for neural networks has a similar role as the shrinkage factor  $\nu$  for boosting.

Model selection for neural networks is performed by keeping track of the validation likelihood over iterations of gradient descent. The model with the best performance on the validation set is returned. Since all conditionals are learned at once only a single model selection is needed. The most analogous selection procedure in our approach is the one based on a linear ordering of the trees for all dimensions (see Section 3.3).

#### 4.4. Likelihood evaluation

A comprehensive list of the computational complexities for likelihood evaluation in neural autoregressive networks was provided by Germain et al. (2015). The simplest autoregressive model based on neural networks is NADE (Larochelle and Murray, 2011), which consists of a single hidden layer. Likelihood evaluations for NADE can be performed with  $\mathcal{O}(DH)$  operations, where  $H$  is the number of hidden units. Deeper versions of NADE (Uria et al., 2014) with multiple layers of hidden units require  $\mathcal{O}(DH^2)$  operations though. Autoencoder networks (Germain et al., 2015) with more than one layer of hidden units require  $\mathcal{O}(DH + H^2 + D^2)$  operations. For ensemble networks (Uria et al., 2014) the complexities scale multiplicatively with the number of different orderings of the data dimensions. Neural networks, which use stochastic hidden units (Gregor et al., 2014; Bornschein and Bengio, 2015), require calculations that are exponential in  $H$  because a sum over all possible hidden configurations has to be computed. In that case approximations via importance sampling are typically used (Salakhutdinov and Murray, 2008). In contrast to that, the computational complexity to evaluate the exact log-likelihood of a test sample under our model is  $\mathcal{O}(DTL)$ , where  $L$  is the average tree depth. If the trees are balanced then  $L = \log J$ . Thus, likelihood evaluations with our approach are orders of magnitude faster compared to the sophisticated variants of neural autoregressive networks.

#### 4.5. Potential benefits of our approach

We summarize what we believe are benefits of our approach over neural autoregressive networks. As discussed in Section 4.1 and 4.4, our model has computational advantages because the training procedure can be parallelized and likelihoods can be calculated efficiently. This makes it possible to scale the approach to very high dimensions.

We see an additional advantage in terms of hyperparameter tuning. The only two hyperparameters that we have to choose in our approach are the shrinkage factor  $\nu$  and the number of tree leaves  $J$  (the number of trees  $T$  is chosen through one of the model selection procedures from Section 3). Choosing  $\nu$  is easy. Indeed, with two exceptions (for speed reasons) we were able to use the same shrinkage factor for all our experiments (Section 5). To choose  $J$  we explored a handful of possible values. In contrast to that, successfully tuning neural networks is still considered an art (Bengio et al., 2013). Besides the (initial) learning rate, the number of gradient descent iterations and the architecture of the neural network there are several “optimization tricks” that might be required in order to achieve satisfactory performance. Typical factors, which affect the convergence, are the mini-batch size, the learning rate decay schedule and the momentum of the gradient descent learning algorithm. Moreover, careful initialization is important and additional regularization procedures like weight decay or dropout (Srivastava et al., 2014) might be needed to generalize well.

The capacity of neural networks is typically chosen such that all potentially useful connections exist. Indeed, the functional structure (2) of the conditional distributions in neural autoregressive networks implies that  $x_d$  always depends on all predictors  $\mathbf{x}_{1:d-1}$ . The hope is that at the time of convergence the connections

to irrelevant variables have a weight close to zero. In contrast to that, our model automatically selects relevant variables because LogitBoost builds the functional structure of the conditional distributions in a sequential manner. Apart from being more interpretable, sparse sets of predictors are well known to perform better for high-dimensional problems (Guyon and Elisseeff, 2003). In Section 5.2 we experimentally show that our model performance is little affected by a large number of irrelevant variables.

A usual step after model selection is refitting of the parameters using the pooled training and validation data. In our model this is easy to perform. For each dimension we keep the learned functional structure fixed, i.e., we use exactly the same trees. We then simply rerun the updates of the leaf values according to equation (4). This refitting is computationally efficient because the most expensive part of training, namely deciding which split to perform, can be skipped. Updating the model parameters with additional data is less natural in neural networks. One heuristic to choose the number of iterations for the combined data is to learn until the average training likelihood from the earlier validated model is reached (Uria et al., 2014). However, this does not really specify a particular model complexity and the approach will not work well if the validation data is easier or harder to fit than the training data. In our model such a retraining would correspond to learning new tree structures, which is a much slower process than simply updating the parameters.

Finally, boosted trees provide a simple way to quantify the importance of dependent variables. It is customary to sum up the gain (3) for each predictor, which corresponds to the reduction in uncertainty due to that variable. In Section 5.2 we exemplify the usefulness of this tool. We are not aware of a comparable feature importance measure for neural networks.

## 5. Experiments

### 5.1. Quantitative performance

We evaluate our LogitBoost autoregressive network (LBARN) on 9 standard benchmark datasets using the same splits into training, validation and test sets as in previous work. The ratio of sample size and dimensionality in these datasets varies between 1 and 300, see Table 1 for details. The MNIST dataset (Salakhutdinov and Murray, 2008) was binarized through random sampling based on the original intensity values. The OCR-letters dataset comes from the Stanford AI Lab<sup>2</sup> and the remaining seven datasets come from the UCI Machine Learning Repository (Lichman, 2013). We first focus on a quantitative comparison with neural autoregressive models. The standard way to assess the quality of probabilistic models is by evaluating the likelihood of the test data under the trained models. In principle, other proper scoring rules (Bickel, 2007) could be considered but there are no previously reported results for alternative scoring rules. Reported likelihoods for the used datasets can be found in Germain et al. (2015) and Bornschein and Bengio (2015).

LBARN was trained using up to  $T = 1,000$  rounds of boosting per dimension. The number of trees to be used for each dimension was chosen through the model selection procedures presented in Section 3 (we used the forward variant of the linearized selection procedure). For the number of leaves per tree we considered  $J \in \{2, 4, 8, 16, 32, 64, 128\}$  (for balanced trees this corresponds to depths of 1-7) and selected the best value based on validation performance. The shrinkage factor was set to  $\nu = 0.02$ . This value turned out to be small enough to train trees with a large number of leaves. Even smaller values of  $\nu$  did not lead to better results but only slowed down the learning process. In cases where convergence had not occurred after 1,000 rounds of boosting we increased  $\nu$  accordingly. This was only necessary for two datasets.

Table 1 reports our obtained test log-likelihoods together with previous results. We also included the selected number of leaves  $J$  for each dataset and the used shrinkage factor  $\nu$ . For datasets with somewhat larger sample size (relative to the data dimension) deeper trees turned out to be beneficial. Table 1 also shows the test log-likelihoods after refitting of the parameters (in the network with individual model selection) using the pooled training and validation data. This always led to a small improvement. When comparing with previous results for fairness we only considered the model without refitting. Overall LBARN is among the

---

<sup>2</sup><http://ai.stanford.edu/~btaskar/ocr/>

Table 1: Average log-likelihoods (in nats) per test example for various models and datasets. The best result (without refitting) for each dataset is shown in bold. Dataset sizes, standard errors (for LBARN with individual model selection) and selected hyperparameters are shown in *italic*.  $J$  is the selected number of leaves per tree and  $\nu$  is the shrinkage factor.

[illegible]



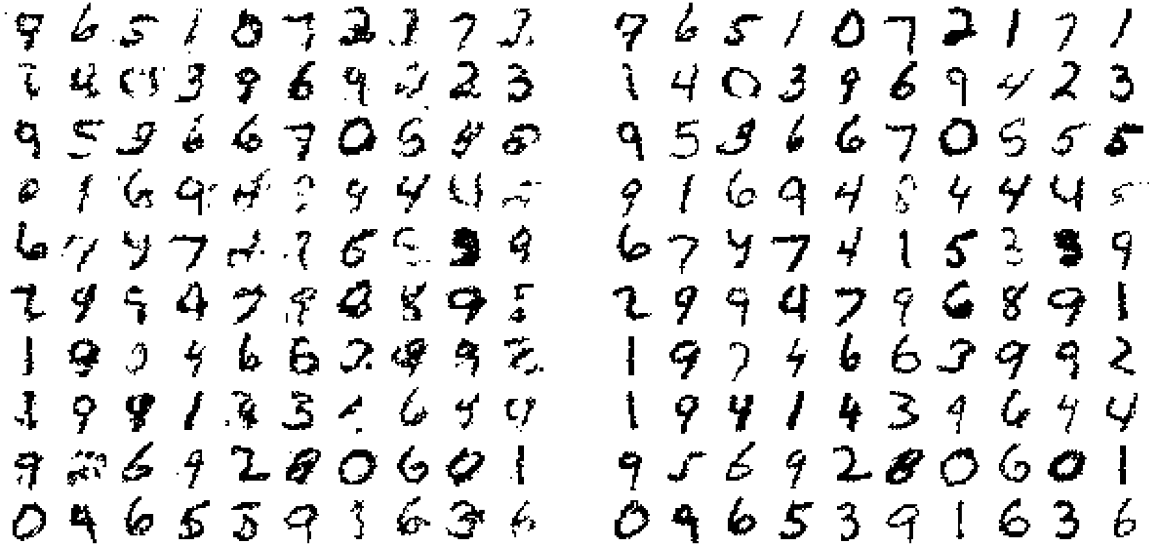


Figure 1: **Left:** Samples from LBARN trained on MNIST digits. **Right:** Closest training examples.

best performing autoregressive models. On 5 datasets our results are better than the best reported result for neural autoregressive networks. On the other 4 datasets the performance is competitive with the state of the art. For the MNIST digits our model performs better than a neural network with a single layer of hidden units (NADE) and as good as a two-layer autoencoder with randomized connectivity pattern (MADE). The models that perform better than LBARN use a large ensemble of models for various orderings of the data dimensions (EoNADE) or rely on stochastic hidden units, which makes exact likelihood computations intractable (DARN, RWS-NADE). In Figure 1 we show samples of the learned model for MNIST digits, together with nearest neighbors from the training set.

As an additional comparison we trained autoregressive networks using a single probability estimation tree (Provost and Domingos, 2003) for each dimension. The number of leaves for each tree was determined based on the performance on the validation set. In addition, small pseudocounts were used to regularize the probabilities. For each dataset the same regularization strength was chosen for all trees using the validation set. The obtained test log-likelihoods of these single-tree autoregressive networks can also be found in Table 1. With one exception boosted trees performed much better than single trees.

The performance of the three model selection procedures is overall comparable. Individual validation for each dimension worked well in most cases. Overfitting of the validation set only occurred for the NIPS-0-12 dataset, which has 500 dimensions but contains merely 100 validation samples. Selection of a common number of trees for all dimensions was often the worst option. The linearized model selection worked almost as well as individual selection and performed better for datasets with a small number of validation samples. The standard errors reported in Table 1 for the test log-likelihoods of LBARN correspond to the network with individual model selection.

Figure 2 shows a map of the selected number of trees at each pixel for the MNIST dataset. For many dimensions around 400 trees were enough. However, a large fraction of the conditionals use close to 1,000 trees, which is the total number of boosting rounds. In those cases the holdout performance saturated. That means more trees do not lead to overfitting but the additional gain is very small. To confirm this experimentally, we continued training for up to 2,000 rounds of boosting. The test log-likelihood increased slightly from the original -86.69 to -86.62.

The main hyperparameter, which we have to choose for LBARN, is the number of leaves  $J$  for each tree. In the left panel of Figure 3 we show the validation log-likelihood on the OCR-letters dataset as a function of boosting rounds for three different choices of  $J$ . The shrinkage factor was always  $\nu = 0.02$ . For  $J = 32$




Figure 2: Map of the selected numbers of trees for the MNIST dataset.




Figure 3: **Left:** Validation log-likelihood on the OCR-letters dataset for different numbers of leaves. **Right:** Variable importance for the central pixel (white) of the OCR-letters data.

LogitBoost did not fully converge after 1,000 rounds of boosting. For  $J = 64$  the best validation performance is obtained with around 600 rounds of boosting, after that slight overfitting begins. For  $J = 128$  strong overfitting occurs after 400 rounds of boosting. Importantly, the maximum of all three curves is almost the same. So, (at least in this example) the exact value of  $J$  is not that important as long as it is in the appropriate range.

A variant of LogitBoost (Friedman et al., 2000) is Gradient Boosting (Friedman, 2001), which approximates the gradient of the Bernoulli log-likelihood to learn boosted trees rather than using a second-order expansion. We ran a few experiments with this alternative boosting method. The performance was usually a bit worse than with LogitBoost. This confirms the observation made by Li (2010) that second-order information improves the results.

### 5.2. Variable importance and variable selection

The splitting gain (3) is a useful tool to quantify variable importance. We illustrate this through the model trained for the OCR-letters data (consisting of images of size  $16 \times 8$ ). In the autoregressive network the central pixel (located at row 8, column 4) is modeled conditional on all pixels that are earlier in the rowwise ordering (i.e. that are above or to the left). As shown in the right panel of Figure 3, the dependencies are mostly local. Indeed, the four nearest neighbors (left, top-left, top and top-right pixel) together explain

about 62% of the uncertainty. Being able to quantify the dependencies between the variables is a useful feature, which improves the understanding of the dataset.

We performed an additional experiment to illustrate that the conditional distributions in LBARN can be learned robustly even when a large number of irrelevant variables is present. The OCR-letters data has 128 dimensions. By stacking three independent copies next to each other we created a new dataset with 384 dimensions, corresponding to the product of three (independent) OCR-letters distributions. Specifically, the training, validation and test datasets are formed by using the original data for dimensions 1-128 and by using two different shufflings of the sample orders for dimensions 129-256 and 257-384, respectively. The best achievable log-likelihood for this dataset is three times the log-likelihood of the OCR-letters data. If during the training phase spurious relationships between the copies are found then the training log-likelihood will be larger than three times the training log-likelihood of the OCR-letters data. However, in that case the test likelihood will be smaller than three times the test log-likelihood of the OCR-letters data. For the first copy  $\mathbf{x}_{1:128}$  the average test log-likelihood we achieved with our trained model was -24.60, as before. In our autoregressive network, the second copy  $\mathbf{x}_{129:256}$  is estimated conditional on the first 128 (irrelevant) variables and the third copy  $\mathbf{x}_{257:384}$  is estimated conditional on 256 nuisance variables. The test log-likelihood for the second (-25.25) and third (-25.42) copy are only a bit lower compared to the first copy. In particular, the performance is still better than the best result from the alternative models in Table 1. A possible reason why the performance decreased at all is that all our trees have the same number of leaves. Consequently, some splitting variables have to be chosen. By requiring a minimum amount of improvement for each additional leaf we believe that the results in the presence of irrelevant variable can be further improved.

### 5.3. Ordering of the variables

For any permutation  $\pi$  of the dimensions  $\{1, \dots, D\}$  the product of conditional distributions

$$\prod_{d=1}^D \mathbb{P}(x_{\pi(d)} | x_{\pi(1)}, \dots, x_{\pi(d-1)})$$

is equal to the joint distribution  $\mathbb{P}(\mathbf{x})$ . However, if the conditionals are learned from data then different decompositions can yield different results. So far we always used the ordering of the variables in which they are stored in the computer. For specific applications alternative orderings can be more appropriate. For example, by using an ordering in which the first few variables provide as much information as possible about the complete data we can create a simple (noisy) compression mechanism. Such an ordering is characterized by the fact that the later conditional distributions have a low entropy. A greedy strategy to find an ordering with that property is to sequentially add the variable with the highest conditional entropy given all previously added variables (Geman and Jedynak, 1996).

We used the training examples from the MNIST digits dataset to sort the  $28 \times 28 = 784$  variables by conditional entropy. The top-left panel of Figure 4 visualizes the obtained ordering. We consider the increasing-entropy order, which starts from background pixels at the boundary of the image grid and circulates towards the inside, as well as the decreasing-entropy order, which starts from central pixels and moves outwards. To illustrate the captured amount of information we plot in the top-right panel of Figure 4 the cumulative log-likelihoods  $\sum_{d=1}^D \log \mathbb{P}(x_{\pi(d)} | x_{\pi(1)}, \dots, x_{\pi(d-1)})$  as a function of  $D$ . When sorted by increasing conditional entropy the variables at the end of the ordering contribute most. When sorted by decreasing conditional entropy most of the contributions to the log-likelihood come from variables at the beginning of the ordering. Nevertheless, the average test log-likelihoods for the full data are rather similar. The original ordering leads to a test log-likelihood (after refitting) of -86.49. With the increasing-entropy order the log-likelihood is slightly lower (-86.53) while with the decreasing-entropy order it is 1.6 standard errors higher (-86.15). The advantage of having the most informative variables early in the ordering is that we can reduce the data to a small fraction of the dimensions and can still reconstruct the full data well. This is illustrated in the bottom panel of Figure 4, where we use examples from the MNIST test set and generate samples conditioned on the first 10% of the variables.

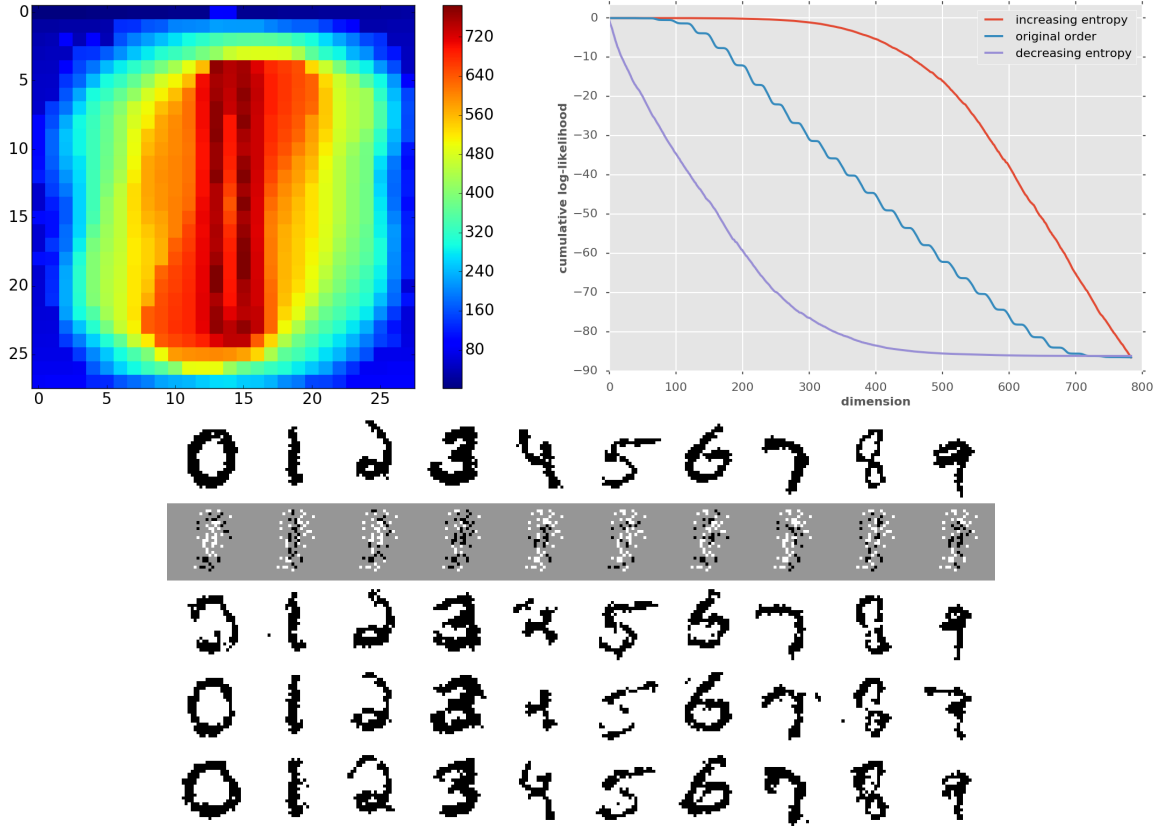


Figure 4: **Top-left:** Pixels for MNIST digits ordered by conditional entropy (blue=low entropy, red=large entropy). **Top-right:** Cumulative log-likelihoods of the MNIST test digits for different pixel orderings. **Bottom:** Compression of MNIST digits. 1st row: Test examples. 2nd row: Reduced data consisting of 78 pixels (gray indicates missing values). 3rd-5th rows: Generated samples conditioned on the reduced data.

## 6. Discussion

We introduced a novel autoregressive network that achieves performance comparable to or even better than the best existing models. Remarkably, this was possible by using separate conditional distribution estimators with no weight sharing. The choice of LogitBoost for modeling the conditionals turned out to be fruitful. However, we conjecture that competitive results can be obtained with other state-of-the-art probability estimators. The specific advantages of our approach compared to neural autoregressive models are simplicity, interpretability and scalability. The Python implementation of LBARN is available at <http://github.com/goessling/lbarn>.

We believe that our model can be further improved by using recently developed boosting extensions like Bayesian additive regression trees (Chipman et al., 2010), dropouts for additive regression trees (Rashmi and Gilad-Bachrach, 2015) or feature subsampling for LogitBoost (Chen and Guestrin, 2016). Moreover, rather than starting the boosting process with log-odds of 0 we could initialize the model by first fitting a logistic regression of  $x_d$  on  $\mathbf{x}_{1:d}$ . The obtained predicted probabilities can then be used as the base probabilities for boosting. Such an initialization potentially decreases the required number of boosting rounds. It also allows for linear effects of the predictors instead of piecewise constant relations. In neural networks this roughly corresponds to introducing direct connections between  $\mathbf{x}_{1:d-1}$  and  $x_d$ .

Weight sharing between conditional distributions for different dimensions or weight sharing between models trained for different orderings of the variables (Uria et al., 2014) is orthogonal to our approach.

However, by replacing the decision trees in our model with directed acyclic graphs (Benbouzid et al., 2012; Shotton et al., 2013) it is in principle possible to introduce weight sharing into our network.

In neural autoregressive networks for real-valued data (Uria et al., 2013) weight sharing among conditional distributions has been shown to be essential for regularization. The natural extension of our approach to real-valued data is to train separate regression estimators for each dimension. This could, for example, be done using L2-boosting (Friedman, 2001; Bühlmann and Yu, 2003), Gaussian process regression (Rasmussen and Williams, 2006) or sparse additive models (Ravikumar et al., 2009). Note that in addition to the conditional mean one also has to estimate the conditional variance, assuming that a simple parametric model for the conditional distributions is used. Alternatively, a multi-class LogitBoost method could be used if the real values are discretized.

## Acknowledgements

I am grateful to Yali Amit for his comments on this work. I also thank the anonymous reviewers for their valuable feedback, which helped to improve the manuscript.

## References

- B. J. Frey, Graphical Models for Machine Learning and Digital Communication, MIT press, 1998.
- G. F. Cooper, E. Herskovits, A Bayesian method for the induction of probabilistic networks from data, *Machine learning* 9 (4) (1992) 309–347.
- M. Goessling, Y. Amit, Mixtures of Sparse Autoregressive Networks, in: *International Conference on Learning Representations (workshop)*, 2016.
- S. Bengio, Y. Bengio, Taking on the curse of dimensionality in joint distributions using neural networks, *IEEE Transactions on Neural Networks* 11 (3) (2000) 550–557.
- H. Larochelle, I. Murray, The neural autoregressive distribution estimator, in: *International Conference on Artificial Intelligence and Statistics*, 29–37, 2011.
- Y. Bengio, Discussion of the neural autoregressive distribution estimator, in: *International Conference on Artificial Intelligence and Statistics*, 38–39, 2011.
- B. Uria, I. Murray, H. Larochelle, A deep and tractable density estimator, in: *International Conference on Machine Learning*, 467–475, 2014.
- T. Raiko, Y. Li, K. Cho, Y. Bengio, Iterative neural autoregressive distribution estimator nade-k, in: *Advances in Neural Information Processing Systems*, 325–333, 2014.
- M. Germain, K. Gregor, I. Murray, H. Larochelle, MADE: Masked autoencoder for distribution estimation, in: *International Conference on Machine Learning*, 881–889, 2015.
- L. Breiman, J. Friedman, C. J. Stone, R. A. Olshen, *Classification and Regression Trees*, CRC press, 1984.
- F. Provost, P. Domingos, Tree induction for probability-based ranking, *Machine Learning* 52 (3) (2003) 199–215.
- C. Boutilier, N. Friedman, M. Goldszmidt, D. Koller, Context-specific independence in Bayesian networks, in: *International Conference on Uncertainty in Artificial Intelligence*, 115–123, 1996.
- N. Friedman, M. Goldszmidt, Learning Bayesian networks with local structure, in: *Learning in graphical models*, Springer, 421–459, 1998.
- J. H. Friedman, Greedy function approximation: A gradient boosting machine, *The Annals of Statistics* (2001) 1189–1232.
- J. Friedman, T. Hastie, R. Tibshirani, Additive logistic regression: A statistical view of boosting, *The Annals of Statistics* 28 (2) (2000) 337–407.
- N. Shafik, G. Tutz, Boosting nonlinear additive autoregressive time series, *Computational Statistics & Data Analysis* 53 (7) (2009) 2453–2464.
- T. J. Hastie, R. J. Tibshirani, *Generalized additive models*, vol. 43, CRC press, 1990.
- P. Li, Robust logitboost and adaptive base class (abc) logitboost, in: *Conference on Uncertainty in Artificial Intelligence*, 302–311, 2010.
- D. Mease, A. J. Wyner, A. Buja, Boosted classification trees and class probability/quantile estimation, *The Journal of Machine Learning Research* 8 (2007) 409–439.
- K. Gregor, I. Danihelka, A. Mnih, C. Blundell, D. Wierstra, Deep autoregressive networks, in: *International Conference on Machine Learning*, 1242–1250, 2014.
- J. Bornschein, Y. Bengio, Reweighted wake-sleep, in: *International Conference on Learning Representations*, 2015.
- R. Salakhutdinov, I. Murray, On the quantitative analysis of deep belief networks, in: *International Conference on Machine Learning*, 872–879, 2008.
- Y. Bengio, A. Courville, P. Vincent, Representation learning: A review and new perspectives, *IEEE transactions on pattern analysis and machine intelligence* 35 (8) (2013) 1798–1828.
- N. Srivastava, G. E. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: a simple way to prevent neural networks from overfitting., *Journal of Machine Learning Research* 15 (1) (2014) 1929–1958.

- I. Guyon, A. Elisseeff, An introduction to variable and feature selection, *Journal of machine learning research* 3 (Mar) (2003) 1157–1182.
- M. Lichman, UCI Machine Learning Repository, URL <http://archive.ics.uci.edu/ml>, 2013.
- J. E. Bickel, Some comparisons among quadratic, spherical, and logarithmic scoring rules, *Decision Analysis* 4 (2) (2007) 49–65.
- D. Geman, B. Jedynak, An active testing model for tracking roads in satellite images, *Pattern Analysis and Machine Intelligence* 18 (1) (1996) 1–14.
- H. A. Chipman, E. I. George, R. E. McCulloch, BART: Bayesian additive regression trees, *The Annals of Applied Statistics* (2010) 266–298.
- K. V. Rashmi, R. Gilad-Bachrach, DART: Dropouts meet Multiple Additive Regression Trees, in: *International Conference on Artificial Intelligence and Statistics*, 489–497, 2015.
- T. Chen, C. Guestrin, Xgboost: A scalable tree boosting system, in: *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794, 2016.
- D. Benbouzid, R. Busa-Fekete, B. Kégl, Fast classification using sparse decision DAGs, in: *International Conference on Machine Learning*, 951–958, 2012.
- J. Shotton, T. Sharp, P. Kohli, S. Nowozin, J. Winn, A. Criminisi, Decision jungles: Compact and rich models for classification, in: *Advances in Neural Information Processing Systems*, 234–242, 2013.
- B. Uria, I. Murray, H. Larochelle, RNADE: The real-valued neural autoregressive density-estimator, in: *Advances in Neural Information Processing Systems*, 2175–2183, 2013.
- P. Bühlmann, B. Yu, Boosting with the L-2 loss: regression and classification, *Journal of the American Statistical Association* 98 (462) (2003) 324–339.
- C. E. Rasmussen, C. K. I. Williams, *Gaussian processes for machine learning*, Adaptive computation and machine learning, MIT Press, 2006.
- P. Ravikumar, J. Lafferty, H. Liu, L. Wasserman, Sparse additive models, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 71 (5) (2009) 1009–1030.